

World Record: Large-Scale Molecular-Dynamics Simulation of 19 Billion Particles


Kai Kadau (T-14), Timothy C. Germann (X-7), and Peter S. Lomdahl (T-11);
kkadau@lanl.gov

In the early 1990s, Lomdahl, Beazley, and coworkers did pioneering work on large-scale molecular-dynamics (MD) simulations [1, 2] with up to 131 million particles on a variety of parallel platforms, including the Connection Machine 5. Ever since, large-scale MD simulations have been used intensively to get insight into atomic processes on the pico-second time scale for various solid-state material science issues. It has been demonstrated that on today's largest parallel supercomputers MD simulations with one billion atoms are feasible.

However, such ultra-large-scale simulations present an analysis problem due to the enormous datasets that are produced. For example, the memory needed to save one configuration of one billion atoms in single precision is more than 15 Gbyte, which makes it problematic and inconvenient to save all the raw data and analyze the data in a postprocessing fashion, even on modern platforms with Tbytes of hard-disc capacity. For example, in the configuration of the (X-, Y-, Z-) particle position each of these four numbers takes usually four bytes in single precision. However, for a restart of the system usually the position and velocity vectors are needed in double precision, which makes the memory consumption about 45 Gbyte for one billion particles. Rather, an on-the-fly analysis during the simulation run is desired.

The MD code SPaSM [1, 2] (Scalable Parallel Short-range Molecular dynamics) has multiple tools to analyze the simulation results during the course of the simulation. The methods implemented include generation of pictures of configurations (particle as well

Figure 1—
Timings for runs on a different number of PN for varying number of atoms. Time used per integration step on a fixed number of processors increases well linearly with the number of atoms for system sizes up to some 20 billion. A slight increase of the time used can be seen, if the number of atoms per PN is fixed and the number of PN is increased. The data is compared to earlier demonstration runs by Roth et al. using 128, 256, 512 PN on a T3E (left to right).



as cell based), local and global quantities like stress tensor elements, temperature, energy, displacement, centrosymmetry parameter, etc. It also allows for calculating histograms and correlation functions for the whole system as well as for subsets of the system (see [3] for examples).

Here, we demonstrate the capability of the SPaSM library to run and analyze systems containing up to approximately 20 billion atoms on the QSC-machine at Los Alamos National Laboratory [4]. Furthermore, we introduce a new parallel sphere rendering capability that we developed in order to generate high-resolution pictures of configurations containing millions to billions of particles. To the best of our knowledge this is the largest MD simulation reported to date. According to the good linear scaling of the simulation time with the number of particles, it would be possible to run 100 billion atoms on a large partition of the current Q-machine. This potential enables future simulation of a physical problem on the micrometer-scale, which so far only has been theoretically investigated by continuum methods.

The flexibility of the SPaSM library can be exploited in other ways, for instance by opening up the possibility of large-scale agent-based simulations. Smaller-scale agent-based models for the spread of epidemics can now scale up to encompass the entire world population of 6.3 billion people.

- [1] P.S. Lomdahl, P. Tamayo, N. Gronbech-Jensen, and D.M. Beazley, *Proceedings of Supercomputing 93*, G.S. Ansell, Ed. (IEEE Computer Society Press, Los Alamitos, CA, 1993), p. 520.
- [2] D.M. Beazley and P.S. Lomdahl, *Comput. Phys.* **11**, 230 (1997).
- [3] K. Kadau, T.C. Germann, P.S. Lomdahl, B.L. Holian, *Science* **296**, 1681 (2002).
- [4] K. Kadau, T.C. Germann, P.S. Lomdahl, *Int. J. Mod. Phys. C* **15**, 193 (2004).




Figure 2—
Approximately 37 million particles rendered on four PN with a resolution of 5000 pixel by 5000 pixel (top). The bottom shows a close-up of the same picture file. The color represents the potential energies of the atoms: dark gray = -6 and red = -2.